DOCUMENT RESUME

ED 105 713 FL 005 728

AUTHOR Schachter, Jacquelyn
TITLE What Must Computer Understanding Systems Understand?
PUB DATE Nov 73
NOTE 22p.; Paper presented at the meeting of the Acoustical Society of America (86th, November 1973)

EDRS PRICE MF-$0.76 HC-$1.58 PLUS POSTAGE
DESCRIPTORS Artificial Languages; *Computational Linguistics; Computer Science; Descriptive Linguistics; English; Grammar; *Programing Languages; *Semantics; Sentence Structure; *Structural Analysis; *Syntax; Vocabulary

ABSTRACT

This paper attempts to characterize some of the knowledge that people have of their language's semantic structure, in order to determine what a computer has to know to perform the operations of searching for facts and answering questions. Necessarily, this sort of semantic analysis is done in connection with syntactic structure. The meaningful elements of English are divided into sets of words and sets of sentences. Their general properties are examined as well as the general relations between elements. The paper concludes with a discussion of the controversy between generative semanticists and interpretivists over syntactic and semantic structure. The work of logicians on artificial language is cited as being helpful in explanations of syntactic and semantic structure in natural languages. (AM)

ED105713

BEST COPY AVAILABLE

WHAT MUST COMPUTER UNDERSTANDING SYSTEMS UNDERSTAND?

by

Jacquelyn Schachter

Paper presented at the 86th meeting of the Acoustical Society of America, November 1973.

FL 005 728

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
NATIONAL INSTITUTE OF EDUCATION
THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

I understand that a number of you are interested in developing computer understanding systems, and more specifically, question answering systems. Let's assume, for an hour, that we all have the same goal in mind. We would like to construct a system that understands some fragment of the English language in written form. The system will understand a limited vocabulary, and a limited number of syntactic constructions. It will accept only sentences that are requests in question form, and sentences which are assertions in declarative form. It won't, for example, be able to accept sentences in imperative form, conditional sentences, etc.

The system will have some facts about certain subjects stored in its memory in some canonical form. It will be able to represent the meaning of a sentence in this canonical form, and will have a way to translate from English into this canonical form. The task of the system is to use its knowledge, in conjunction with a problem-solving system, to understand certain questions posed to it, interpret those questions as specific commands, search its memory for the relevant facts, and produce the appropriate answers.

This is the goal. The question now becomes, what would the system have to know in order to perform the operations involved in understanding the question, searching for the facts, and giving the answer. Another way of stating this question is to ask what knowledge the speakers of English possess that allows them to perform the same operations.

You will notice that many examples in this paper are stated in terms of a fact, a question, and an answer. Let's assume that the facts represent knowledge that our system has. The questions represent requests for information that we have made of the

system, and the answers represent responses we would like the system to make. Given these facts and these questions, we can focus on what the system would have to know in order to come up with the responses listed. Now, if we were operating with an actual system, only the questions and the answers would be represented in English. The facts would be stored in some notation which was independent of actual English words. But of course there is no notation available, that we're all familiar with, so I have been forced to represent the facts in English also. I have tried to represent the facts in such a way that they are semantically transparent. I may have failed in some cases, and I apologize in advance for that.

My job today, then, is to try to characterize some of the knowledge of the semantic structure of their language that people have. I think I can do that, but I can't do it in isolation from syntactic structure. I don't think that one can draw sharp distinctions between syntax and semantics.

Most people who study languages assume, correctly I think, that the meaning of a sentence is composed of the meanings of its parts. In order to understand the semantic structure of a sentence we must know what its parts are and how these parts are related to each other. In descriptive terms we can say that these parts are related to each other. In descriptive terms we can say that:

A. a description of the syntactic structure of a language will tell us what the elements in the language are, and how they can be combined.

B. a description of the semantic structure will tell us what the meanings of the elements are, and what the meanings of the combinations are. To

attempt to do semantics without syntax would be to describe the meanings of the elements and their combinations without knowing what the elements and their combinations are.

Because of time limitations, I would like to divide the meaningful elements of English into two groups - the set of words, and the set of sentences. This, of course, is a gross over-simplification which obscures many semantic facts that need to be characterized. But by limiting myself to only two groups in this way, I can better focus in the general properties that they exhibit and the general relations that they enter into.

Normally we use words to describe and to refer to what's going on in the world. The word Russian, in the phrase the Russian is, in general, used to pick out some particular individual with certain characteristics. And the word hit, as in the sentence the Russian hit the ball, is used to pick out some particular activity of a certain type. But not all words have these functions of describing and referring in all cases. These is a small group of words which can serve as syntactic indicators with little if any descriptive content, or referential use.

1. It annoyed the Russian that he couldn't hit the ball

In (1) both it and that have no descriptive content, and they don't refer to anything. It is functioning as a sort of place for the subject position and that serves to mark the following clause as subordinate.

If you take a descriptive word like boy, you would want to say that it has meaning which is independent of the object you use it to refer to. You might want to characterize it in terms of

semantic primes like 'human', 'male','young'. But there is a class of words, called deictic words or indexicals, which have little content out side of a particular context. That is, we use them to refer but not to describe.

2. a) There's a Russian who's in my class

b) There's a Russian who's in my lab who doesn't like Egyptians

In (2b) the word there is most commonly used as a syntactic indicator. It doesn't have meaning and doesn't refer. In (2a) the word there, and the word my, are both used to refer to something. But without knowledge of the situation in which this sentence occurs, we don't know what they refer to, and we can't say anything about whether the sentence is true or not. The implications for a question answering system are obvious. The system will have to make decisions about when a word doesn't refer, and when it does. And if the word refers, the system will have to decide what, in the present context, the word refers to.

3. fact: Three people are in the lab: a Russian, an Egyptian, and Henry

Q: Is there a Russian in the room with me?

A: Yes

In order to answer this question, the system has to know that the word there does not refer, and that the word me does refer, and also that the referent of the word me is not the Russian. The referent must be either the Egyptian or Henry.

Another more complicated case is the use of one in the following example:

4. fact: One country left the UN and one country joined NATO

Q: Did one country leave the UN and join NATO?

A: No

The phrase one country tells us that we are referring to a country, in each case, but not which country. We need a more complete context to decide that. But in addition, the fact that the phrase one country is mentioned twice, in the fact, tells us that two different countries are being referred to. The Q asks if one country did two things. The fact tells that two countries did one thing each. To recapitulate what I've said so far - we can describe words in terms of their descriptive properties, and their referential properties. Words can be said to refer, and they can be said not to refer. Of the words that are used to refer, some can be said to have descriptive content independent of context, others have only context dependent content.

But even for words which always have descriptive content, it is not enough to know their meanings in isolation. They enter into relations with other words which add to their meaning. We sometimes call this added meaning categorical meaning, as opposed to cognitive meaning.

5. a) The Russian told the Egyptian

b) The Russian fell into the Suez Canal

c) The Russian looked nervous

In each case we know the meaning of the word Russian, and in each case the word is the subject of the sentence, and yet we perceive that the relation of the word to the verb is different in each case. In (a), we know the Russian is the actor; in (b), he is the experiencer or patient; in (c), he is neither. A question answering system has to know this kind of information in order to

decide what, of all the information it has stored about the Russian, will be appropriate as an answer to different kinds of questions. Let's assume that we have (a), (b), (c) stored as facts in the system. In answer to question 6, only (a) is an appropriate answer.

6. Q: What did the Russian do?

A: The Russian told the Egyptian

In answer to 7, only (b) is appropriate.

7. Q: What happened to the Russian?

A: He fell into the Suez Canal

There is no generalized question form which will get at the information in (c). In order to get at that kind of information, the question would have to be much more specific.

Words also enter into relations with the syntactic structures in which they occur. This interrelation serves to delimit the meaningful functional role which the word has in that structure. An example of the relations between words and the structure they occur in is the following.

8. fact: The Russian promised the Egyptian (to go).

Q: Did the Russian promise to go?

A: Yes

9. The Russian expected (the Egyptian to go)

Q: Did the Russian expect to go?

A: No

In this case, it is the difference in the underlying structural relation that the two verbs enter into which determine that the Russian would do the going in 8, and that the Egyptian would do the going in 9. The parentheses in the facts in 8 and 9 are there to show that the information is not stored in an ambiguous syntactic

form in the system. But the questions asked are expressed in the same syntactic form, and it is the relation of the verb to the contruction which will disambiguate these cases.

I'd like to leave the discussion of words now, and go on to a discussion of sentences. Many properties that we can predicate of words, we can also predicate of sentences.

Let us first focus on properties of sentences, and then go into a discussion of the relations that sentences can enter into. Ther are a number of properties that sentences can have. They can be analytic, they can be contradictory, they can be ambiguous. The first two properties, although they might be of some interest to logicians will not be, I think, of interest to those interested in our Q answering system, so I will ignore them. But the property of ambiguity of sentences is one that is essential for us to face. A sentence can be ambiguous in different ways; it can be structurally ambiguous because one surface structure can represent two or more underlying structures; it can be lexically ambiguous because one or more words have different meanings, which have not been eliminated by the meanings of the other words in the sentence; and a sentence can be referentially ambiguous - what I have in mind here is in fact that we can use language to refer to objects or events in the world, but we can also use it to refer to the descriptive content of the words in the sentence itself.

Structural ambiguity occurs when different underlying structures are altered in such a way as to make them look the same when they are written. If a certain lexical item can occur in two different underlying structures, and both underlying structures have the same surface structure, then that resulting surface structure is ambiguous. Look first at the question of 10. It is ambiguous. But in both interpretations, the meanings of the words are the same. It's the difference in syntactic relations which gives us the

difference in semantic interpretation.

10. fact: The members of the Russian delegation are not dignitaries.

fact: The members of the Russian delegation are visiting several dignitaries.

Q: Are the members of the Russian delegation visiting dignitaries?

A: What do you mean by "visiting dignitaries"?

Normally a speaker uses previous information to disambiguate this question, and one approach would be to program our system to attempt to do the same thing. It might look through its store of information on the members of the Russian delegation to see how they are characterized. If the machine knows that the Russians are dignitaries who are visiting, it will answer "yes"; if it has the fact that the R's are visiting important people, it will answer "yes" also. But this tactic doesn't always work for people, and it won't work for the system either. Suppose the system knows the two facts stated in 10. Here the system would have to disambiguate. The correct answer to one interpretation of the question would be "Yes - they are visiting several dignitaries" and the answer to the other interpretation would be "No - the Russians are not dignitaries."

Lexical ambiguity occurs when we use one word to refer to different objects or events or relations in the world. Most words in the language are ambiguous in this way; there are only a few, usually highly technical words that have one and only one meaning. Normally we can eliminate all but one meaing of a word in a sentence by referring to the meanings of the other words in the sentence and figuring out what the appropriate combinations are. But

occassionally we have to ask what a particular word means.

11. Q: Did the Egyptians miss the Suez Canal?

A: I don't know what you mean by the word "miss"

Do you mean 1. long for, 2. fail to find?

Referential ambiguity have proven to be far more difficult to handle than the kinds we've talked about so far. Yet it is one that all speakers of a language can comprehend, and one which I'm very much afraid any good question answering system will have to cope with. I mentioned before, when talking about words, that they are used to describe and to refer to things or events, etc. Let me get a bit more explicit here. Words or phrases or sentences have sense (or descriptive content). Words, phrases, sentences are used to refer to objects, events, states of affairs, etc. Let us say of a simple singular noun or noun phrase that its sense is a concept of an individual, and that its referent is an individual. Let us say of a predicate phrase, that its sense is the property which it expresses, and that it is used to refer to a class of individuals (eg. the predicate "is red" refers to the class of things that are red, and has as it's sense the property of being red). Let us say of a sentence that its sense is the proposition expressed by that S and that its used to refer to a state of affairs, or as some say, to a truth value.

So far, I've said that we use words and sentences to refer to objects and classes of objects, and states of affairs. But what is important to be aware of is the fact that we can also use them to refer to their senses; i.e. we use a N or NP to refer not to an individual, but to the individual concept. We can use a predicate to refer not to a class of objects, but to a certain property that the class has, and we can use a sentence to refer not to an actual

state of affairs, but rather to the proposition expressed by the sentence.

12. a) The Premier of Russia went to Egypt yesterday.

b) The Premier of Russia must be elected by the party.

In a), I'm using the term "the Premier of Russia" to refer to an individual. But in the most obvious interpretation of the b) sentence, I'm not using the term "the Premier of Russia" to refer to an individual at all. I'm really saying that any Premier of Russia, whoever he is, must be elected by the party. That is, I'm using the term "the Premier of Russia" to refer to its sense. Of course, in the other interpreation of b), I'm really referring to an individual; I might be implying that that individual is no different than anyone else. So 12b) is ambiguous - in one interpretation, the less likely one in this case, I'm actually referring to an individual; in the other interpretation, I'm not referring to an individual at all. 12a) is not ambiguous in this way. The term "The Premier of Russia" can only have the "referring to an individual" use.

Let us look at 13b) next:

13. a) The Russian does not have a hat.

b) The Russian's hats are always brown.

13b) is a sentence in which in the most obvious interpretation , the predicate, are always brown, is used to refer to its sense rather than to a class of objects. In the other reading, the one in which the predicate is used to refer to a particular class of brown hats, the interpretation would be something like - the Russian's hats are brown and never change to any other color; they stay brown. But notice that if the system had both a) and b) stored as facts, the only reading of b) would be the one in which the terms refer

to their senses, not to objects in the world. Not let us look at a case where a whole clause is ambiguous in this way.

14. a) The Russians don't know the author of the UN speech.

b) The Russians believe that the author of the UN speech is a fool.

In 14b), the clause following believe is ambiguous as to what it refers to. In one interpretation, it refers to a state of affairs. There is a author and he is a fool, and the Russians believe it. But again, if we have 14a) as a fact in the system, the only possible interpretation of 14b) is the one in which the sentence following believe refers to a proposition and not to a state of affairs. 14b) could be true even if there were no author and no UN speech.

This ambiguity in reference is determined by the occurrence in these sentences and others like them, of several verb, modal and adverb classes: propositional attitude Vs like say, believe, think; emotion verbs like want hope, would hate, adverbs like always, necessarily, possibly, modals like must,would. This distinction is one that the philosophers have more to say about than linguists. Until recently, it has been ignored by linguists and you won't find much in linguistic literature on it. Life would be a lot easier if we could ignore it, but I don't think we can. We've seen how it might affect the storage of facts in our system. Let's see how it might affect its question answering ability.

15. fact: The Russians found a solution

Q: What solution did the Russians find?

A: I don't know/the actual solution

16. fact: the Russians sought a solution

Q: what solution did the Russians seek

A: That's a nonsensical question

Now let us turn to relations between sentences. If a speaker of a language knows the meaning of a sentence within that language, he also knows that there are certain relationships that hold between that sentence and certain other sentences, and he knows what these other sentences are. In order for a person to be able to answer questions about facts, he must understand these relations between sentences. Our question answering system must do this also. Probably the most obvious relation between sentences, one that everyone would expect me to mention, is synonymy. So I'll mention it - synonymy. But I don't want to spend any time on it because it's notoriously difficult to define, and I think maybe we can get away with not defining it directly, both for our question answering system, and for a semantic theory of natural language. I think a lot of synonymy relations between sentences can be accounted for in other ways - for example, the relation between the sentences "It took me a long time to write this paper" and "I took a long time to write this paper" can be accounted for in terms of their both having the same underlying structure. The relation that I want to focus on here is what I'm calling entailment. In general, entailment is a relation between two sentences such that if one sentence is true, then the other sentence has to be true also. But entailment takes several forms, and it is instructive to look at them in some detail.

17. a) All of the dignitaries are Russian

b) Some of the dignitaries are Russian

If we know a) to be a fact, we also know b) to be a fact because we know the relationship between <u>all</u> and <u>some</u>. We've known about this kind of entailment for a long time - anyone who's had a logic course will be aware of it. But we have been finding that this

entailment relation holds between many more types of sentence pairs than had been previously noticed.

18. a) The Russians managed to warn the Egyptians.

b) The Russians forgot to warn the Egyptians.

c) The Russians expected to warn the Egyptians.

If we know that a) is a fact, we also know that the Russians warned the Egyptians. Another case of entailment. On the other hand, if we know that b) is a fact, we know that the Russians didn't warn the Egyptians - a case of negative entailment. But even if we know that c) is a fact, we do not know if the Russians did or did not warn the Egyptians. Verbs like manage and forget are called implicative verbs.

There is a special kind of entailment which is sometimes called the "converse relation." This is a relation of mutual entailment; i.e., if A entails B, then B entails A.

19. fact: Russia sold arms to Egypt.

Q: Did Egypt buy arms from Russia?

A: Yes

We know that if Russia sold arms to Egypt, then Egypt bought arms from Russia. And if Egypt bought arms from Russia, then Russia sold arms to Egypt. It makes no difference which is its fact and which is the question. The answer is the same.

There are some interesting entailment relations that show clearly the interrelations between the words and the structural forms with which they occur. Let us look at one involving verbs and tense forms:

20. fact: The Russian began walking in the park today.

Q: did the Russian walk in the park today?

A: yes

21. fact: The Russian began walking to Washington today

Q: did the Russian walk to Washington today?

A: I don't know

The entailment relation between the past progressive and simple past holds in the walk case, but it does not hold for the walk to case.

A sentence can entail many other sentences. But in the usual case what a sentence entails is not what the negation of that sentence entails. The sentence "All of the dignitaries are Russian" entails the sentences "Some of the eignitaries are Russian." But its negation, "All of the dignitaries are not Russian" does not entail "Some of the dignitaries are Russian." But for certain sentence types, both the sentence and its negation entail the same sentence. When this occurs, we call the relation one of presupposition. In general, then for a sentence A to presuppose a sentence B, then both A, and the negation of A must entail B.

22. fact: The Russians (don't) realize that an insurgent group is operating in Egypt.

Q: Is an insurgent group operating in Egypt?

A: yes

Notice that the yes answer is correct for either fact; that the Russians do realize that an insurgent group is operating in Egypt, and that the Russians don't realize that an insurgent group is operating in Egypt. This presupposition relation occurs in this case because the sentence contains a factive verb realize. But it also holds between sentences in which no factive verb appears. The sentence "It was Brutus who killed Caesar," presupposes the sentence "Someone killed Caesar."

The presupposition relation has special significance for

question-answering systems, because questions carry presuppositions too. In order to know the correct answer to a question, you must know what its presuppositions are. Suppose the system knows the fact in 23, but it doesn't know anything about when this occurred.

23. fact: The Russians armed the Egyptians

Q: When did the Russians arm the Egyptians?

A: I don't know.

The question presupposes that the Russians armed the Egyptians and asks for additional information about that fact. Since the system doesn't have the additional information, it correctly answers "I don't know." But suppose that the machine had no information at all about the Russian's arming the Egyptians. If you ask it the same question, "When did the Russians arm the Egyptians", the "I don't know" answer would be misleading. One approach to han-- dling is by having the system reply, "I can't explain a nonexistent fact." This looks like a neat solution. But the problem is even more complicated than this. Suppose the presupposition of the question flatly contradicts a fact that the system has stored, as in 24.

24. fact: the Russians did not arm the Egyptians

Q: when did the Russians arm the Egyptians

A: but they didn't

In this case, the "I don't know" answer would be inadequate, and so would the "non-existent fact" answer. The best answer would be the one which contradicted the presupposition of the question, and therefore made the question irrelevant.

One area that strikes me as being an extremely difficult one for designers of question-answering systems is the one involving the process of logical inference. As speakers of a language, we can distinguish

between valid and invalid arguments expressed in the language. Given two premises, we can tell whether the conclusion follows. One might argue that this is not really a problem of meaning, its a problem of logic. But the decision as to whether a certain conclusion follows from a certain set of premises is dependent on the meanings of those sentences - the meanings of the words, and the structures in which they occur. It would be desirable for our system to be able to make inferences, like the one in 25.

25. fact: The Russians left when the Egyptians arrived.
fact: The Egyptians arrived at 2
Q: When did the Russians leave?
A: when the Egyptians arrived
A: at 2

The first answer is certainly acceptable, and the system does not have to make any inference in this case. The second answer is also acceptable, and the system would have to make an inference. One could argue that this ability to infer need not be programmed into the system because you could get both answers anyway, by making it a two-step question.

Q: When did the Russians leave?
A: When the Egyptians arrived.
Q: When did the Egyptians arrive
A: at 2

But this two-step question process will not always work.

26. fact: All the Egyptians in Moscow are from Cairo.
fact: Mr. Al Faiz is an Egyptian in Moscow.
Q: Is Mr. Al Faiz from Cairo?
A: yes

Given the facts of 26, the system that couldn't make inferences could

only answer "I don't know." This, it seems to me, would be a waste. After all, the system has the information stored. What is lacking is a means to get at that information.

This problem of "getting at" the information stored in the system is closely related to what some have called the "Principle of Cooperation" between speakers, in an area that's usually called pragmatics, not semantics. Assuming two speakers have good will toward each other, they cooperate in such a way that even if the question asked by speaker A cannot be answered precisely by speaker B, speaker B doesn't give up. He cooperates by giving what information he has. Clearly the most useful question-answering system would cooperate in the same way.

27. fact: The Russians warned the Egyptians four times in March, 1972.

Q: How many times did the Russians warn the Egyptians in 1972?

A: I don't know.

A: (at least) four times.

The question is not one for which the system is prepared. But an "I don't know" answer would seem inappropriate. Equally, inappropriate is the answer"four times" - it overgeneralizes. In both cases, the answer would be violating the principle of cooperation. The best answer in this case is "at least four times."

You may have noticed that as I have been talking about presupposition and logical inference and the principle of cooperation I have been saying things like "the best answer" and"the most appropriate answer would be" - that is, I have been giving value judgments. This is because when I think about these areas, I am

less sure, and considerably more speculative, about how closely want our system to mirror the abilities of humans. It may be that for many purposes we could set our sights much lower, and not have to concern ourselves with these areas at all. But I think that we won't know beforehand. We won't know, one way or another, until we build these systems -- and that's the job I leave for you.

Before I finish, I would like to say something about what's going on in linguistic theory these days, with regard to the study of semantic structure of language. I mentioned that one assumption that most of us make is that the meaning of a sentence is composed of the meanings of its parts; that the notion of the part/whole relation is essentially a syntactic one. Many transformationalists have taken this to mean that a syntactic description is logically prior to a semantic description. In fact, in the early days of transformational theory, the syntactic component of a description of a language was seen as not only prior to but also independent of the semantic description of the language. But this assumption has not been held by everyone who is interested in language. Many philosophers, for example, have also been studying languages, both natural and artificial. Many of these philosophers are logicians and logicians are free to construct or describe whatever languages they want. They are not limited as linguists are, to describe some pre-existing languages. When they do construct languages, they construct them in such a way that they part/whole relation given by the syntax is exactly the right part/whole relation needed for the semantics. Linguists in the past have tended to ignore the work done by logicians for exactly this reason. They have argued that they are tied to the facts of natural language as they exist, and not as they would like them to be. What the linguists failed to see,

for a long time, is that the semantic distinctions that logicians wanted to make in their artificial languages were drawn from semantic distinctions that exist in natural language. When linguists began to take semantics seriously, the part/whole problem reared its ugly head again. The argument raging among transformational linguists nowadays is essentially whether the correct part/whole distinctions needed for syntax are the correct part/whole distinctions needed for semantics. Another way of putting this problem is to ask the question "Is the underlying syntactic structure of a language semantically transparent?" One camp, the generative semanticists argue that it is, and the other camp, the interpretivists argue that it isn't. A number of linguists have been turning to the logicians to see what use they can make of their recent work in answering the questions that have been raised. And there are now a number of linguists and logicians who are pursuing the notion of a logic for natural language, "natural logic," one that combines the syntactic insights of the linguists and the semantic insights of the logicians. Anyway, the controversy rages, and it is often distressing to those of us in the middle, although it is probably productive for the field as a whole But I think that in the end, the theoreticians will not be able to satisfactorily answer these questions unless they get feedback from the people who try to put these theories to use. A linguistic theory will be considered good only insofar as it can explain part of what we do during the process of communication. The attempt to simulate this process using the various theories can give valuable aid in helping to decide which theories are better, and for what purposes.